gMark: Schema-Driven Generation of Graphs and Queries

Guillaume Bagan, Angela Bonifati, Radu Ciucanu, George H. L. Fletcher, Aurélien Lemay, Nicky Advokaat

Abstract—Massive graph data sets are pervasive in contemporary application domains. Hence, graph database systems are becoming increasingly important. In the experimental study of these systems, it is vital that the research community has shared solutions for the generation of database instances and query workloads having predictable and controllable properties. In this paper, we present the design and engineering principles of gMark, a domain- and query language-independent graph instance and query workload generator. A core contribution of gMark is its ability to target and control the diversity of properties of both the generated instances and the generated workloads coupled to these instances. Further novelties include support for regular path queries, a fundamental graph query paradigm, and schema-driven selectivity estimation of queries, a key feature in controlling workload chokepoints. We illustrate the flexibility and practical usability of gMark by showcasing the framework's capabilities in generating high quality graphs and workloads, and its ability to encode user-defined schemas across a variety of application domains.

Index Terms—Graph databases, Selectivity estimation, Recursive queries, Benchmarking.

arXiv:1511.08386v5 [cs.DB] 7 Oct 2016

1 Introduction

The problem. We study the problem of schema-driven generation of synthetic graph instances and corresponding query workloads for use in experimental analysis of graph database systems. Our study is motivated by the ubiquity of graph data in modern application domains, such as social and biological networks and geographic databases, to name a few. In response to these pressures, systems that can handle massive graph-structured data sets are under intense active research and development. These systems span from pure graph database systems to more focused knowledge representation systems. Native graph databases such as Neo4j [1] and Sparksee [2] propose their own declarative data model and query language, with particular attention to query optimization, and space and performance. In contrast to this trend of specialized systems, general-purpose systems such as LogicBlox [3] rely on declarative solutions that can cover a broader range of use cases. Furthermore, knowledge representation systems such as Virtuoso [4] implement the standard RDF graph data model and SPARQL query language to handle complex navigational and recursive queries on large-scale Semantic Web data.

Keeping pace with these developments, synthetic graph and query generation and benchmarking solutions for graph data management systems has been a proliferating activity, which started within the Semantic Web community (e.g., [5], [6]) and recently within the database community, e.g., through the extensive work of the LDBC Council [7]. The latter has been the first to propose a chokepoint-driven design of graph database benchmarks, which allows fine-grained control of chokepoints of queries and all aspects of the involved data. By relying on a fixed schema and a carefully designed set of benchmark queries, this lets the community focus on crucial features of query optimization and/or parallel processing, and decreases the confusion and the incomparable results of evaluated systems.

Our solution. We propose a complementary approach in which the focus is not on the design of individual queries but rather on whole query workloads. This approach relies on the control of diversity of both graph schemas and query workloads, which lets us vary the structural properties of data as well as tailor the generated queries to a particular domain or application. We emphasize that a workload-centric approach primarily targets different benchmarking and experimental scenarios from the query-centric approach of current benchmarks. Indeed, we believe that this approach is important in contexts in which multiple queries (i.e., those belonging to a query workload) need to be considered altogether, which occurs in many cases, such as in multi-query optimization, workload-driven database tuning, streaming applications, mapping discovery, and query rewriting in data integration systems.

In this paper, we present gMark, a fully-functional practical framework that realizes this workload-centric perspective. gMark takes a schema-driven approach to the flexible and tightly-controlled generation of synthetic graph instances coupled with sophisticated query workloads. The complete framework is provided as *open-source software*[1] ready for use in the graph processing community.

We further note that the core of query workload generation in gMark revolves around a novel method we propose here for selectivity estimation for graph queries, which is, to the best of our knowledge, the first of its kind. This method is of independent interest, applicable to other contexts such

- *G. Bagan and A. Bonifati are with the University of Lyon 1, Bat. Nautibus, 23-25 av. Pierre de Coubertin, 69622 Villeurbanne Cedex, France.* {guillaume.bagan,angela.bonifati}@liris.cnrs.fr
- *R. Ciucanu is with the University Blaise Pascal, Clermont-Ferrand, 1 rue de la Chebarde, 63178 Aubière Cedex, France. Research partially done at the University of Oxford, supported by EPSRC platform grant DBOnto.* ciucanu@isima.fr
- *G. H. L. Fletcher and N. Advokaat are with TU Eindhoven, P.O. Box 513, 5600 MB Eindhoven, The Netherlands.* {g.h.l.fletcher@,n.advokaat@student.}tue.nl
- *A. Lemay is with the University of Lille 3 & Inria, Parc Scientifique de la Haute Borne, 40 av. Halley, 59650 Villeneuve d'Ascq, France.* aurelien.lemay@inria.fr

1. https://github.com/graphMark/gmark

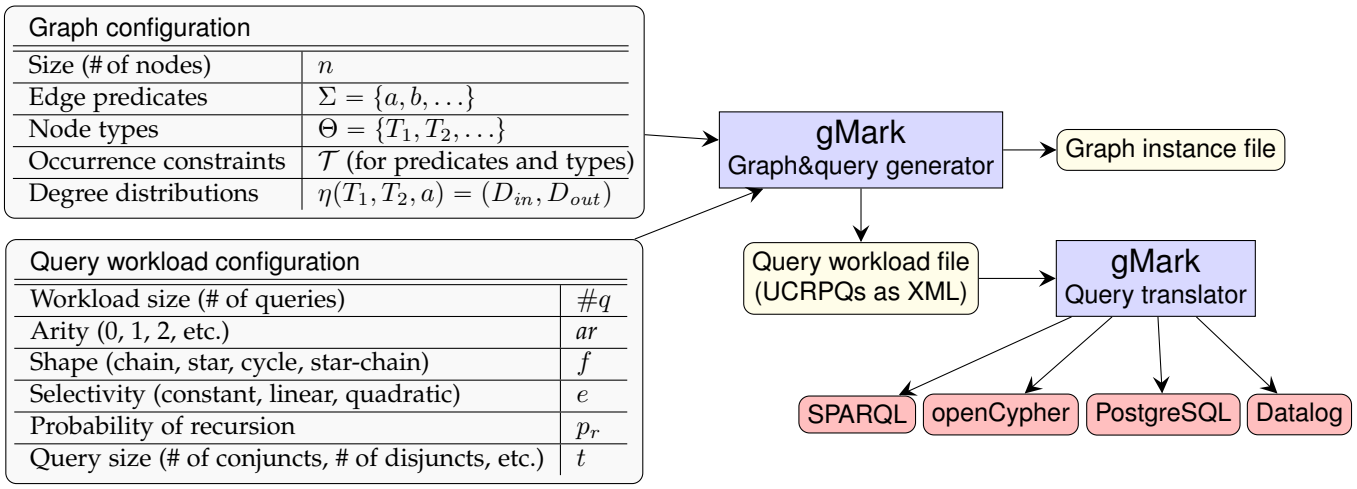
Graph configuration
Size (# of nodes) n
Edge predicates $\Sigma = \{a, b, \ldots\}$
Node types $\Theta = \{T_1, T_2, \ldots\}$
Occurrence constraints $\mathcal{T}$ (for predicates and types)
Degree distributions $\eta(T_1, T_2, a) = (D_{in}, D_{out})$
Query workload configuration
Workload size (# of queries) $\#q$
Arity (0, 1, 2, etc.) ar
Shape (chain, star, cycle, star-chain) f
Selectivity (constant, linear, quadratic) e
Probability of recursion p_r
Query size (# of conjuncts, # of disjuncts, etc.) t
gMark
Graph&query generator
Graph instance file
Query workload file
(UCRPQs as XML)
gMark
Query translator
SPARQL
openCypher
PostgreSQL
Datalog

Fig. 1. Overview of the gMark workflow.

as query optimization and query inference on graphs and, more generally, relational data and queries.

1.1 gMark design principles

To support a broad range of systems and domains, we aim in the gMark design to cover features and capabilities commonly found in graph query processing, graph analytics, and schema validation. Reflecting this, the interplay between the following key features characterizes the gMark system architecture. To the best of our knowledge, gMark is the first framework to satisfy all of these criteria.

Built-in support for schema definition. A major goal of gMark is to account for an array of fundamental user-defined schema constraints during graph generation. In general, gMark is domain-independent, while it can be used to target a rich variety of realistic domains. Hence, graph instance generation leverages an optional schema definition, called a *graph configuration*, which includes the enumeration of predicates (i.e., edge labels) and node types (i.e., node labels) occurring in the data, along with their properties in generated instances (see Figure 1).

Controlled instance and workload diversity. Given a graph configuration, a subsequent challenge is to exploit it for query workload generation. Current approaches in query generation rely on the graph instances to generate queries with desired behavior. However, this approach is unfeasible for large and loosely structured networks. We argue that query workload generation must primarily rely on the graph configuration rather than on the generated graph instances, while still enforcing the desired behavior of the generated queries. Towards this, gMark supports a broad range of schema-driven parameters in both the graph and query workload generator (see Figure 1), e.g., capable of gauging both navigational and recursive query execution performance, while not relying on a fixed set of query templates. To the best of our knowledge, existing solutions for graph databases do not support such configurable (and easily extensible) range of parameters.

Language- and system-independence. Another important design principle of gMark is independence from particular query language syntaxes or systems. Towards this, gMark supports various practical output formats for the graphs and for the queries, including N-triples for data, and SQL, SPARQL, Datalog and openCypher as concrete query language syntaxes for query workloads (see Figure 1). gMark is also easily extensible to support other output formats.

Broad applicability. Finally, gMark aims to broadly support a wide range of application domains. As an example, we show below that it is easy to adapt the scenarios (modulo incomparable features) of three existing state-of-the-art benchmarks into meaningful gMark configurations, while also adding new gMark features: the LDBC Social Network Benchmark [7], SP2Bench [6], and the Waterloo SPARQL Diversity Test Suite (WatDiv) [5].

1.2 Contributions and organization

We depict an overview of the gMark workflow in Figure 1. The goal of this paper is to present the design, engineering, and first empirical study of the gMark framework. We next outline our main contributions.

- We formalize the problems of graph generation and query workload generation, and show that they are intractable in general (Section 3).
- We provide an in-depth presentation of the gMark design principles, for the generation of both graphs (Section 4) and query workloads (Section 5). The most notable novel features are support for recursive queries and query selectivity estimation in the generated query workloads.
- We empirically show the capability of gMark to cover diverse graphs and query workloads, the accuracy of the estimated selectivities, and the scalability of the generator (Section 6).
- We present an in-depth experimental comparison of a representative selection of state-of-the-art graph query engines using gMark, which brings to light important limitations of current graph query processing engines, in particular w.r.t. recursive query processing (Section 7).

We close this section by noting that gMark supports the full range of data and query features, and practical query syntaxes discussed above. Query selectivity tuning is supported only on binary queries (i.e., queries of arity two). This should not be considered as a limitation since

already selectivity tuning is a non-trivial problem for binary queries, and such queries already make interesting major practical cases of query benchmarking in graph databases. For example, all regular path queries (which appear as "property paths" in SPARQL 1.1) are binary.

2 Related work

Data generation and benchmarking frameworks have played an important role in database systems research over the last decades, where efforts such as the TPC Benchmarks and XML benchmarking suites have been crucial for advancing the state of the art [8], [9]. Similarly, in support of the experimental study of graph data management solutions, a variety of synthetic graph tools such as SP2Bench [6], LDBC [7], LUBM [10], BSBM [11], Grr [12] and WatDiv [5] have been developed in the research community. Complementary to the gMark approach, application-driven derivation of graph configurations and synthetic graph instances has also recently been studied [13]–[15]. Furthermore, extensive collections of real-world graphs such as SNAP [16] and KONECT [17] are now available as community resources.

Currently available resources (i) rely on fixed graphs or instances of fixed graph schemas, or (ii) provide limited or no support for generating tailored query workloads to accompany graph instances. These aspects are difficult to jointly relax, especially in the context of loosely structured complex networks. Indeed, there is no community consensus on schema formalisms for graph data, an area that is still in an early stage of investigation [18]. Furthermore, constructing workloads with given selectivity and structural features is a very difficult problem [19]–[25]. As mentioned in the Introduction, current approaches such as WatDiv and LDBC perform selectivity estimation on generated graph instances, which becomes unfeasible when dealing with massive graphs and query workloads.

In gMark, we address this challenge by generating tailored workloads directly from the graph schema, where query selectivity is set as one of the input parameters. While BSBM and WatDiv do support a workload-centric approach, they do not provide this fine-grained level of control of query behavior. In general, we are not aware of any solutions for controlling selectivity during query generation relying solely on graph schemas.

Finally, to our knowledge gMark is the first solution for generating workloads exhibiting recursive path queries. In particular, the queries generated by gMark are the so-called *unions of conjunctive regular path queries* [26]. This fundamental query language covers many graph queries that appear in practice. In particular, SPARQL 1.1 and openCypher have conjunctive regular path queries as their core constructs. They are also expressible in modern Datalog-like query languages [3] and in SQL:1999. As discussed in the Introduction, gMark supports the output of query workloads in all these concrete query language syntaxes.

3 The generation problem

We start this section by intuitively introducing the generation problem via a real-world motivating example that also emphasizes some of the limitations of existing benchmarks (Section 3.1). Then, we formalize the benchmark generation problem i.e., generating a graph instance and a query workload on this instance according to a given set of constraints. More precisely, we formally define the problems of *graph generation* (Section 3.2) and *query workload generation* (Section 3.3), and we detail constraint parameters for both problems. We conclude the section by showing that the generation problem is *intractable* (Section 3.4).

3.1 Motivating example

Assume that a user wants to perform an extensive empirical evaluation of a new graph query processing algorithm that she designed. For this purpose, the user needs to efficiently generate: (i) graphs of different characteristics and sizes (to test the robustness and scalability of her algorithm), and (ii) query workloads sufficiently diverse to highlight strong or weak points of her new development. Additionally, our user would like to specify all parameters in a declarative way and to be able to simulate real-world scenarios.

For instance, the user would like to generate graphs simulating a bibliographical database that uses a simple schema consisting of 5 node types and 4 edge predicates. Intuitively, the database consists of *researchers* who *author papers* that are *published* in *conferences* (held in *cities*) and that can be *extended* to *journals*. Moreover, the user would like to specify constraints on the number of occurrences for both the node types and edge predicates, either as proportions of the total size of the graph or as fixed numbers e.g., as in Figure 2(a) and 2(b). For instance, for graphs of arbitrary size, half of the nodes should be authors, but a fixed number of nodes should be cities where conferences are held (in a realistic scenario the number of authors increases over time, whereas the number of cities remains more or less constant).

Moreover, our user wants to specify real-world relationships between types and predicates via schema constraints e.g., as in Figure 2(c). For instance, the first line encodes that the number of authors on papers follows a Gaussian distribution (the *in-distribution* of the schema constraint), whereas the number of papers authored by a researcher follows a Zipfian (power-law) distribution (the *out-distribution* of the schema constraint). The following lines in Figure 2(c) encode constraints such as: a paper is published in exactly one conference, a paper can be extended or not to a journal, a conference is held in exactly one city, the number of conferences per city follows a Zipfian distribution, etc.

Whereas specifying all aforementioned constraints as an input gMark *graph configuration* (cf. Figure 1) can be easily done via a few lines of XML, to the best of our knowledge there is no benchmark where they can be specified. For instance, in SP2Bench [6] (which is also based on a similar bibliographical scenario), all constraints are hardcoded and the only parameter that a user can specify is the size of the graph, which makes it impossible for the user to finely tune schema-related characteristics of the graph. Moreover, in WatDiv [5], although the user can specify similar global constraints on the node types and the out-distributions, the absence of global constraints on the edge predicates and the absence of in-distributions entail important limitations, such as the absence of control on the selectivities of the queries of the generated query workloads.

In gMark we allow the user to finely tune the selectivities of the generated queries. For instance, the user can

Node type	*Constr.*
researcher	50%
paper	30%
journal	10%
conference	10%
city	100 (fixed)

(a) Node types.

Edge predicate	*Constr.*
authors	50%
publishedIn	30%
heldIn	10%
extendedTo	10%

(b) Edge predicates.

source type $\xrightarrow{\textit{predicate}}$ *target type*	*In-distr.*	*Out-distr.*
researcher $\xrightarrow{\text{authors}}$ paper	Gaussian	Zipfian
paper $\xrightarrow{\text{publishedIn}}$ conference	Gaussian	Unif. [1,1]
paper $\xrightarrow{\text{extendedTo}}$ journal	Gaussian	Unif. [0,1]
conference $\xrightarrow{\text{heldIn}}$ city	Zipfian	Unif. [1,1]

(c) In- and out-degree distributions.

Fig. 2. The bibliographical motivating example.

specify that she wants queries that, for any graph size, have constant, linear, or quadratic selectivity (we formally define these selectivity classes later on in the paper). To the best of our knowledge, no graph database benchmark supports such a feature. In particular, SP2Bench uses a fixed set of queries, while WatDiv can generate synthetic workloads, but without schema-driven selectivity control. As another remarkable difference to the state-of-the-art benchmarks, none of them supports recursive queries such as $(\texttt{authors}\cdot\texttt{authors}^-)^*$ which selects all pairs of researchers linked by a co-authorship path (by $^-$ we denote the predicate inverse and by * the transitive closure). As shown in the input gMark *query workload configuration* in Figure 1, the user can finely-tune e.g., the structure, size, selectivity of such queries.

3.2 Graph generation

gMark generates directed edge-labeled graphs and outputs them in formats that are compatible with the supported query languages. In this section, we formally define a *graph configuration* (cf. Figure 1), which is essentially a set of constraints that generated graph instances should satisfy. We start by giving a definition of the *schema* constraints which are the backbone of graph configurations.

Definition 3.1 *A* graph schema *is a tuple* $\mathcal{S} = (\Sigma, \Theta, \mathcal{T}, \eta)$ *where* Σ *is a finite alphabet of* predicates, Θ *is a finite set of* types *such that each node of the generated graph is associated with exactly one type,* $\mathcal{T}$ *is a set of constraints on* Σ *and* Θ *associating to each predicate and type either a proportion of its occurrences or a fixed constant value, and* η *is a partial function associating to a triple consisting of a pair of input and output types* T_1, T_2 *in* Θ *and a symbol* a *in* Σ*, a pair* (D_{in}, D_{out}) *of in- and out-degree distributions.*

Predicates correspond to *edge labels*, and in the remainder we use the two terms interchangeably. A degree distribution is a *probability distribution*, among which gMark supports *uniform*, *Gaussian* (also known as normal), and *Zipfian* distributions. For each distribution, the user can specify the relevant parameters (i.e., min and max for uniform, μ and σ for Gaussian, and s for Zipfian). If the user wants to specify only the in- or the out-distribution, she can mark the other one as *nonspecified*. Notice that the parameters for the in- and out-degree distributions of each triple T_1, T_2, a have to be consistent in order to guarantee the compatibility of the number of generated ingoing and outgoing edges. We discuss the details of this consistency check in Section 4.

Definition 3.2 *A* graph configuration *is a tuple* $\mathbb{G} = (n, \mathcal{S})$*, where* n *is the number of nodes of the graph and* $\mathcal{S}$ *is the schema of the graph.*

Example 3.3 Take a graph configuration $\mathbb{G} = (n, \mathcal{S})$ s. t.:

- The graph should have $n = 5$ nodes.
- The graph should satisfy the schema $\mathcal{S} = (\Sigma, \Theta, \mathcal{T}, \eta)$, where $\Sigma = \{a, b\}$, $\Theta = \{T_1, T_2, T_3\}$, $\mathcal{T}$ is defined as $\mathcal{T}(T_1) = 60\%$, $\mathcal{T}(T_2) = 20\%$ and $\mathcal{T}(T_3) = 1$ and η is defined as follows (we report only some constraints):

$$\eta(T_1, T_1, a) = (g, z), \quad \eta(T_1, T_2, b) = (u, g),$$
$$\eta(T_2, T_2, b) = (g, ns), \quad \eta(T_2, T_3, b) = (ns, u),$$

where by $u, g, z,$ and ns we denote uniform, Gaussian, Zipfian, and non-specified distributions, respectively. For

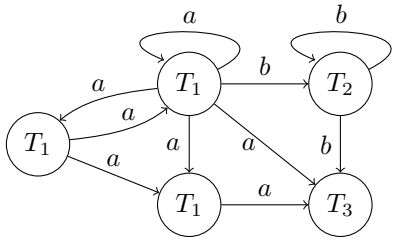

Fig. 3. Graph from Example 3.3.

instance, the graph in Figure 3 can be generated by using this graph configuration. Although a much larger graph is needed to observe the actual distributions, we refer to Section 5 for further examples also handling distributions. ▫

3.3 Query workload generation

We next formally define *query workload configurations* (cf. Figure 1). Towards this, we first outline the query language supported by gMark. As motivated in Section 2, we focus on generating *unions of conjunctions of regular path queries* (UCRPQ). This fundamental query language covers many queries which appear in practice, including the core constructs of SPARQL 1.1 queries, Neo4j's Cypher queries, and many Datalog-based encodings [26].

Recall that Σ is a finite alphabet (cf. Definition 3.1) and let $\Sigma^+ = \{a, a^- \mid a \in \Sigma\}$, where a^- denotes the *inverse* of the edge label a. Let $V = \{?x, ?y, \ldots\}$ be a set of variables and $n > 0$. A *query rule* is an expression of the form

$$(\overline{?v}) \leftarrow (?x_1, r_1, ?y_1), \ldots, (?x_n, r_n, ?y_n)$$

where: for each $1 \leqslant i \leqslant n$, it is the case that $?x_i, ?y_i \in V$; $\overline{?v}$ is a vector of zero or more of these variables, the length of which is called the *arity* of the rule; and, for each $1 \leqslant i \leqslant n$, it is the case that r_i is a regular expression over Σ^+ using $\{\cdot, +, *\}$ (i.e., *concatenation*, *disjunction*, and *Kleene star*). Without loss of generality, we restrict regular expressions to only use recursion (i.e., the Kleene star symbol $*$) at the outermost level. Hence, expressions can always be written to take either the form $(P_1 + \cdots + P_k)$ or the form $(P_1 + \cdots + P_k)^*$, for some $k > 0$, where each P_i is a *path expression* i.e., a concatenation of zero or more symbols in Σ^+. We refer to the right-hand side of a query rule as the *body* of the query rule, each subgoal $(?x_i, r_i, ?y_i)$ of the body as a *conjunct*, and to the left-hand side as the *head* of the query rule.

A *query* $Q \in$ UCRPQ is a finite non-empty set of query rules, each of the same arity. The semantics $Q(G)$ of evaluating Q on a given graph G (having edge labels in Σ) is the standard one following that of unions of conjunctive Datalog queries [26], assuming standard *set-oriented* semantics. In summary, a query is basically a collection of *query rules*, each rule having several *conjuncts*, each conjunct having several *disjuncts* whose paths have a certain *length*.

Example 3.4 Take the following UCRPQ query:

$$(?x, ?y, ?z) \leftarrow (?x, (a \cdot b + c)^*, ?y), (?y, a, ?w), (?w, b^-, ?z)$$
$$(?x, ?y, ?z) \leftarrow (?x, (a \cdot b + c)^*, ?y), (?y, a, ?z)$$

This query selects nodes x, y, z such that one can navigate between x and y with a path in the language of $(a \cdot b + c)^*$, and moreover, can navigate between y and z with a path in the language of $a \cdot b^- + a$. This query consists of two rules consisting of three conjuncts and two conjuncts, resp. The conjuncts of the form $(?x, (a \cdot b + c)^*, ?y)$ have two disjuncts (of length 2 and 1, resp.) and all other conjuncts have only one disjunct (of length 1). □

We define *query size* as a tuple

$$t = ([r_{\min}, r_{\max}], [c_{\min}, c_{\max}], [d_{\min}, d_{\max}], [l_{\min}, l_{\max}])$$

providing intervals of minimal and maximal values for the number of rules, conjuncts, disjuncts, and length of the paths in the query, resp., that generated queries should have. For example, the query from Example 3.4 has size $([2, 2], [2, 3], [1, 2], [1, 2])$. In gMark, users can specify minimal and maximal values for all of these parameters; in turn, the query generation algorithm can assign values that range in these intervals. For simplicity of presentation, we assume in the remainder that a query consists of only one rule.

Definition 3.5 *A* query workload configuration *is a tuple* $\mathbb{Q} = (\mathbb{G}, \#q, ar, f, e, p_r, t)$ *where* $\mathbb{G}$ *is a graph configuration,* $\#q$ *is the number of queries in the workload (defined on all instances of* $\mathbb{G}$*),* ar *is the arity constraint,* f *is the shape constraint,* e *is the selectivity of the queries in the workload,* p_r *is the probability of recursion, and* t *is the query size.*

Notice that in addition to the graph configuration $\mathbb{G}$ (cf. Section 3.2), the user can specify several other constraints. First, ar is the range of allowed arities for the queries in the workload. For instance, the query from Example 3.4 has arity 3. We also support Boolean queries (arity 0). The shape constraint f contains the supported query shapes (among which chain, star, cycle, and star-chain are supported in gMark) and the user can specify which among them she would like to have in the generated query workload. Similarly, the selectivity constraint e contains the desired selectivity classes, among which we support constant, linear and quadratic (cf. Section 5.2). The user can further specify the probability to have the multiplicity * above a disjunct, reflected by the parameter p_r.

We finally point out that gMark is able to *translate* the generated UCRPQ in four concrete syntaxes (cf. Figure 1): SPARQL, openCypher, PostgreSQL, and Datalog.

3.4 Intractability of the generation problem

In this section, we prove the *intractability* of the problems of graph and query workload generation. First, we prove that the graph generation problem is intractable.

Theorem 3.6 *Given a graph configuration* $\mathbb{G}$*, deciding whether there exists a graph satisfying* $\mathbb{G}$ *is NP-complete.*

Prior to presenting the proof of Theorem 3.6, we would like to introduce some standard macros for encoding pairs of in- and out-distributions. Precisely, we use:

- "1" for non-specified in-degree distribution and uniform out-degree distribution with min=max=1. In other words, $\eta(T_1, T_2, a) = 1$ means that from a node of type T_1 there is precisely one outgoing a-labeled edge to a node of type T_2, and that in a node of type T_2 we can have an arbitrary number of incoming a-labeled edges from nodes of type T_1.
- "?" for non-specified in-degree distribution and uniform out-degree distribution with min=0 and max=1.
- "0" for non-specified in-degree distribution and uniform out-degree distribution with min=max=0.

Proof We show the NP-hardness by reduction from the $\mathrm{SAT}_{\text{1-in-3}}$ problem, known to be NP-complete [27]. Take a 3CNF formula $\varphi = C_1 \wedge \ldots \wedge C_k$ over variables $x_1, \ldots, x_n$. We construct a schema $\mathcal{S}_\varphi = (\Sigma_\varphi, \Theta_\varphi, \mathcal{T}_\varphi, \eta_\varphi)$ and a subsequent graph configuration $\mathbb{G}_\varphi = (n_\varphi, \mathcal{S}_\varphi)$ as follows:

- n_φ: The graph should have $2 \times n + k + 1$ nodes.
- Σ_φ: There should be $3 \times n + k$ symbols (predicates) in the alphabet: $\Sigma_\varphi = \{c_1, \ldots, c_k, b_1, \ldots, b_n, t_1, f_1, \ldots, t_n, f_n\}$
- Θ_φ: There should be $3 \times n + k + 1$ node types: $\Theta_\varphi = \{A, C_1, \ldots, C_k, B_1, \ldots, B_n, T_1, F_1, \ldots, T_n, F_n\}$.
- $\mathcal{T}_\varphi$: There should be precisely one node of type A in the graph, which can be expressed as $\mathcal{T}_\varphi(A) = 1$. Additionally, $\mathcal{T}_\varphi(B_1) = \ldots = \mathcal{T}_\varphi(B_n) = \mathcal{T}_\varphi(C_1) = \ldots = \mathcal{T}_\varphi(C_k) = 1$. Notice that all these constraints could be alternatively expressed by saying that the proportion of each of them should be $1/(2 \times n + k + 1)$ of the graph nodes.
- η_φ such that:
  - $\eta_\varphi(A, T_1, t_1) = \ldots = \eta_\varphi(A, T_n, t_n) =?$
  - $\eta_\varphi(A, F_1, f_1) = \ldots = \eta_\varphi(A, F_n, f_n) =?$
  - $\eta_\varphi(T_i, C_{l_1}, c_{l_1}) = \ldots = \eta_\varphi(T_i, C_{l_m}, c_{l_m}) = \eta_\varphi(T_i, B_i, b_i) = 1$, for every $i \in \{1, \ldots, n\}$, where $c_{l_1}, \ldots, c_{l_m}$ correspond to the clauses in which the variable x_i appears in a positive literal.
  - $\eta_\varphi(F_i, C_{l_1}, c_{l_1}) = \ldots = \eta_\varphi(F_i, C_{l_m}, c_{l_m}) = \eta_\varphi(F_i, B_i, b_i) = 1$, for every $i \in \{1, \ldots, n\}$, where $c_{l_1}, \ldots, c_{l_m}$ correspond to the clauses in which the variable x_i appears in a negative literal.
  - $\eta_\varphi(X, Y, b) = 0$ for all other combinations of types X, Y and symbols b not present in one of the aforementioned cases.

We illustrate the construction for $\varphi_0 = (x_1 \vee \neg x_2 \vee x_3) \wedge (\neg x_1 \vee x_3 \vee \neg x_4)$, for which we obtain (we omit all cases with 0):

- $\eta_\varphi(A, T_1, t_1) = \eta_\varphi(A, F_1, f_1) = \ldots = \eta_\varphi(A, T_4, t_4) =$
  $= \eta_\varphi(A, F_4, f_4) =?$
- $\eta_\varphi(T_1, C_1, c_1) = \eta_\varphi(T_1, B_1, b_1) = \eta_\varphi(F_1, C_2, c_2) =$
  $= \eta_\varphi(F_1, B_1, b_1) = \eta_\varphi(T_2, B_2, b_2) = \eta_\varphi(F_2, C_1, c_1) =$
  $= \eta_\varphi(F_2, B_2, b_2) = \eta_\varphi(T_3, C_1, c_1) = \eta_\varphi(T_3, C_2, c_2) =$
  $= \eta_\varphi(T_3, B_3, b_3) = \eta_\varphi(F_3, B_3, b_3) = \eta_\varphi(T_4, B_4, b_4) =$

$= \eta_\varphi(F_4, C_2, c_2) = \eta_\varphi(F_4, B_4, b_4) = 1$

We claim that $\varphi \in \mathrm{SAT}_{1\text{-in-}3}$ iff there exists a graph that satisfies the graph configuration.

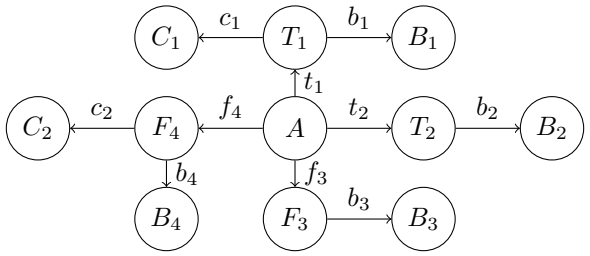

Fig. 4. Graph from the proof of Theorem 3.6.

For the *only if* part, take a valuation that satisfies exactly one literal of each clause and construct a graph that encodes this valuation, starting from a node of type A. For example, for the above formula φ_0 and the valuation such that x_1 and x_2 are true, and x_3 and x_4 are false, construct the graph from Figure 4. Since we choose exactly one among T_i and F_i for every $i \in \{1, \ldots, n\}$ we have exactly one B_i. Moreover, since exactly one literal of each clause is satisfied we have exactly one C_i (for $i \in \{1, \ldots, k\}$). Thus, the constraints in $\mathcal{T}_\varphi$ are satisfied. As for the number of nodes, we have $2 \times n$ (because there is exactly one valuation of each variable and we also have its corresponding B_i) + k (because of the k clauses) + 1 (the node of type A). Consequently, the constructed graph satisfies all the constraints from the configuration.

For the *if* part, take a graph satisfying the constraints. Since it satisfies the constraints from $\mathcal{T}_\varphi$, the graph should have one node A, one B_i (for $i \in \{1, \ldots, n\}$) and one C_i (for $i \in \{1, \ldots, k\}$). Since the total size of the graph is of $2 \times n + k + 1$ nodes and seen how we can reach B_i's and C_i's based on the schema, we infer that the other n nodes correspond to nodes of type T_i or F_i (encoding a valuation of the variable x_i). Since we have precisely one B_i, we infer that each variable has exactly one valuation and since we have precisely one C_i, we infer that in each clause there is exactly one literal that is satisfied. This means that the formula φ is in $\mathrm{SAT}_{1\text{-in-}3}$.

To show the membership of the problem to NP, we point out that a non-deterministic Turing machine has to guess a graph having as many nodes as the constraint from the configuration. The size of such a graph is thus polynomial in the size of the input and testing whether it satisfies the schema can be easily done in polynomial time. □

As a natural consequence of Theorem 3.6, we have that the query workload generation problem is also intractable. That means that some parameters of the query workload cannot be fulfilled and it is not possible to test this efficiently. Hence, gMark follows a heuristic strategy in the generation. More precisely, it attempts to achieve the exact values of the parameters and it may decide to relax some of them in order to obtain linear running time. However, we do not claim that our algorithm is the best linear time algorithm. Nonetheless, in the experimental study we observed fast generation time and good quality graph instances and queries.

Corollary 3.7 *Given a query workload configuration $\mathbb{Q}$, deciding whether there exists a query workload satisfying $\mathbb{Q}$ is NP-complete.*

Proof Recall that a query workload configuration $\mathbb{Q}$ is a tuple $(\mathbb{G}, \#q, ar, f, e, p_r, t)$. Since the graph configuration $\mathbb{G}$ is part of the input of $\mathbb{Q}$, we can take for $\mathbb{G}$ precisely the same encoding of $\mathrm{SAT}_{1\text{-in-}3}$ as in the proof of Theorem 3.6 and then arbitrary values for the other constraints in $\mathbb{Q}$. Then, both *if* and *only if* parts for the NP-hardness follow precisely as in the proof of Theorem 3.6. As for the membership of the problem to NP, we point out that a non-deterministic Turing machine has to guess a query workload and then we can easily decide in polynomial time whether it satisfies the input constraints. In particular, the query workload is of polynomial size (since the size is given as parameter in the query workload configuration), and the graphs that we need to check whether the query workload satisfies the selectivity constraints are also of polynomial size (since the graph size is given as part of the graph configuration that is also part of the query workload configuration). □

Despite the intractability of the generation problems and the gMark heuristic approach, we would like to already point out that the gMark graph and query generator leads to highly accurate results, as we detail with our experiments in Section 6.2, where we report on the selectivity estimation of the generated queries over the generated graphs.

4 Graph generation

As shown in Theorem 3.6, given a graph configuration $\mathbb{G}$, deciding whether there exists a graph satisfying $\mathbb{G}$ is *NP-complete*. Consequently, it is not always possible to generate in PTIME a graph satisfying a given graph configuration $\mathbb{G}$. Nevertheless, generating graph instances in PTIME is an important goal in designing a graph generator, essentially because of *scalability*. To achieve scalability we must therefore relax some constraints specified in the graph configuration.

For all such reasons, the gMark *generation algorithm* (Algorithm 1) takes a heuristic approach that guarantees *linear time* generation (i.e. linear in the size of the input and output). The algorithm pseudocode, which is conceptually quite straightforward, is illustrated below. For each constraint $\eta(T_1, T_2, a)$, it generates a set of edges as follows. It creates a vector v_{src} for the source nodes containing each of the nodes of type T_1 a number of times drawn according to the out-distribution D_{out}. It initializes (line 2), fills (lines 3-4), and, respectively, shuffles (line 7) this vector. The length n_{T_1} of this vector is computed either as $n \times \mathcal{T}(T_1)$ if $\mathcal{T}(T_1)$ is a proportion, or simply $\mathcal{T}(T_1)$ if $\mathcal{T}(T_1)$ is a fixed value. The algorithm creates a vector for the target nodes v_{trg} in a similar manner (lines 2, 5–7). Next, it simultaneously iterates over the two vectors v_{src} and v_{trg} and outputs a number of a-labeled edges corresponding to the minimal length of the two vectors (lines 8-9). The function $id_T(j)$ then returns the j^{th} node of type T in the graph (assuming that the graph nodes are uniquely identified). We note that the random draws in Lines 4 and 6 are statistically independent and hence the order in which we process j has no impact. Likewise, the processing of constraints in Line 1 is statistically independent and hence order-independent.

Algorithm 1 Graph generation algorithm.

Input: A graph configuration $\mathbb{G} = (n, \mathcal{S} = (\Sigma, \Theta, \mathcal{T}, \eta))$
Output: A set of edges (source node, label, target node)

1: **foreach** $\eta(T_1, T_2, a) = (D_{in}, D_{out})$ **do**
2: **let** v_{src} and v_{trg} be two empty vectors
3: **for** $j \in \{1, \ldots, n_{T_1}\}$ **do**
4: **add** $draw(D_{out})$ occurrences of j in v_{src}
5: **for** $j \in \{1, \ldots, n_{T_2}\}$ **do**
6: **add** $draw(D_{in})$ occurrences of j in v_{trg}
7: *shuffle*(v_{src}); *shuffle*(v_{trg})
8: **for** $i \in \{1, \ldots, \min(|v_{src}|, |v_{trg}|)\}$ **do**
9: **output** $(id_{T_1}(v_{src}[i]), a, id_{T_2}(v_{trg}[i]))$

Algorithm 1 produces an output graph *in linear time*, without backtracking. Our implementation leverages further optimizations, notably exploiting the average information of the Gaussian distributions to avoid entirely constructing the vectors when the in- and/or out-distributions are Gaussian. Our experiments in Section 6.2 confirm the practical efficiency and scalability of our graph generator.

Due to the heuristic nature of our algorithm, some of the constraints in the graph configuration might not be fulfilled in the output graph. In particular, for each constraint it generates a number of edges equal to the minimal length of the two vectors (line 8). Consequently, whenever the two vectors have different sizes, then the generated graph may contain nodes that do not satisfy the precise values dictated by the in- or out-distributions. Notice that these distributions are essential in the proof of Theorem 3.6 for showing the intractability of the generation problem.

Due to practical reasons, we decided that our generator should always return a graph to the user instead of performing a potentially costly satisfiability check and possibly aborting the generation. Moreover, this choice still allows to globally preserve the types of distribution for each constraint $\eta(T_1, T_2, a)$, even though the generated number of edges may not satisfy the exact parameters of the in- or out-distributions. In fact, our method relies on the types of distributions (uniform, Gaussian, Zipfian) and not on the actual parameters of the distributions. Such a choice of preserving the global distributions in the graph generator also turns to be useful in the query generator, where we adopt a novel technique of selectivity estimation for the generated queries. As shown in the experiments in Section 6.2, the obtained selectivities exhibit high accuracy, thus confirming the effectiveness of our method.

5 Query Generation

A core innovation of gMark is that graph instances and query workloads are both generated from graph schemas. This allows the tight coupling of queries to instances while still also supporting fine-grained control of the diversity of query workloads. As discussed in Section 3.3, query generation in gMark is guided by the same input schema used for graph generation, which makes the queries of the workload pertinent to graph instances.

In Section 5.1, we present our *query generation algorithm*, which ensures query diversity in terms of *arity*, *shape* (chain, star, cycle, star-chain), *recursion* (probability to have Kleene star above conjuncts), and *size* (number of conjuncts, disjuncts, length of paths). We have already implemented in gMark this algorithm supporting the diversity of all the aforementioned parameters.

Algorithm 2 Query generation algorithm.

Input: A query workload configuration $\mathbb{Q} = (\mathbb{G} = (n, \mathcal{S}), \#q, ar, f, e, p_r, t)$
Output: A set of queries

1: **for** $i \in \{1, \ldots, \#q\}$ **do**
2: **let** $skeleton = get\_query\_skeleton(f, t)$
3: $add\_projection\_variables(skeleton, ar)$
4: **let** $Q = instantiate\_placeholders(skeleton, \mathcal{S}, p_r, t)$
5: **output** Q

Additionally, users can specify constraints on the query *selectivity*. In gMark, we guarantee the selectivity estimation for binary queries, a natural and broad class in the context of graph queries, strictly containing the regular path queries (i.e., property paths in SPARQL 1.1). Selectivity estimation for such queries already involved the development of novel sophisticated machinery. Indeed, in Section 5.2 we formalize the notion of *selectivity classes for binary queries* and we explain the additional data structures and techniques needed to accommodate selectivity estimation for binary queries in the general algorithm. The schema-driven techniques that we introduce for selectivity estimation are non-trivial and, to the best of our knowledge, novel. Moreover, binary queries already make interesting practical cases for graph database benchmarking. In particular, our study based on such queries points out important limitations of existing graph database engines, as we highlight in Section 7.

5.1 General algorithm

The *query generation algorithm* (Algorithm 2) takes as input a *query workload configuration* $\mathbb{Q}$ (cf. Section 3.3) and outputs a set of queries.

For each query, we first generate a query skeleton for the body of the query (line 2) by considering the query shape and the number of conjuncts. A query skeleton is a set of conjuncts of the form $(?x_1, P, ?x_2)$ where $?x_1$ and $?x_2$ are variables, and P is a placeholder that we instantiate afterwards with a regular expression. For example, if we want a chain query and a number of conjuncts $c_{\min} = c_{\max} = 3$, we obtain a query skeleton $(?x_1, P_1, ?x_2), (?x_2, P_2, ?x_3), (?x_3, P_3, ?x_4)$.

The other query shapes are essentially combinations of different *chain queries*. More precisely, we implemented *cycle queries* as two chain queries that share the same endpoint variables, *star queries* as combinations of chains that have the same starting variable, and *star-chain queries* as combinations of chain queries and star queries.

Next, we randomly pick a set of projection variables such that their number is consistent with the arity constraint (line 3); in gMark, we already support the generation of queries of arbitrary arity (including Boolean i.e., arity 0). We finally instantiate the placeholders (line 4) with regular expressions that satisfy the probability of recursion and the other aspects of the size (number of disjuncts and length of paths), and we output the obtained query (line 5).

Algorithm 2 works in linear time in the size of the input and output. It produces a query workload coupled to the graph schema, and satisfying the arity, shape, recursion and size constraints. The Algorithm 2 presented here disregards the selectivity constraint e that is the trickier parameter of a graph configuration and treated in the next section.

5.2 Selectivity estimation algorithm

In this section, we present an innovative method to estimate the selectivity of binary queries, that we already implemented in gMark. Our method impacts the step from line 4 in Algorithm 2, and requires new developments that we introduce next.

In Section 5.2.1, we classify the binary queries in three classes depending on the asymptotic behavior of the size of their result. In Section 5.2.2, we propose a method to estimate the selectivity. In Section 5.2.3, we introduce the data structures needed to implement it. In Section 5.2.4, we revisit Algorithm 2 to accommodate selectivity estimation.

5.2.1 Constant, linear, and quadratic queries

The *selectivity* of a query Q on a graph G is the number of results returned by the evaluation of Q on G, i.e., $|Q(G)|$. This number depends on both the topology and the actual size of the graph instance, given that instances of different sizes can be generated upon the same input schema.

Given a binary query Q and a schema $\mathcal{S}$, for every graph G satisfying $\mathcal{S}$, we assume that the value $|Q(G)|$ behaves asymptotically as a function of the form $|Q(G)| = \beta|G|^{\alpha}$, where α and β are real constants[2]. We say that the above value α is the *selectivity value of* Q (*w.r.t.* $\mathcal{S}$), denoted $\alpha_{\mathcal{S}}(Q)$ or simply $\alpha(Q)$ when it does not lead to ambiguity. Thus, the selectivity value of a query is by definition bounded by the query arity. Since we focus on binary queries, we consider selectivity values such that $0 \leqslant \alpha(Q) \leqslant 2$.

We identify three practical query classes, depending on whether $\alpha(Q)$ is closer to 0, 1, or 2:

• *Constant queries* (for which $\alpha(Q) \approx 0$) select a number of results that does not grow (or barely grows) with the graph size. For instance, a query selecting pairs (`country`, `language`) is constant if the graphs follow a realistic schema specifying that the numbers of countries and languages do not grow with the graph size, and hence the number of query results is more or less constant.

• *Linear queries* (for which $\alpha(Q) \approx 1$) select a number of results that grows at a rate close to the growth of the number of nodes in the graph instances. For example, a query selecting pairs (`language`, `user`) is linear if the schema specifies that the number of users grows with the graph, whereas the number of languages is more or less constant. Another example of a linear query is (`user`, `address`) if we assume that the schema specifies that each user has precisely one address and the number of users grows linearly with the graph.

• *Quadratic queries* (for which $\alpha(Q) \approx 2$) select a number of results that grows at a rate close to the growth of the square of the number of nodes in the graph instances. For example, the transitive closure of the `knows` predicate in a social network is quadratic because a realistic schema should specify that this predicate follows a power-law (e.g., Zipfian) in- and out-distribution. Thus, the query result contains Cartesian products of subsets of users that know and are known by some hub users of the social network.

2. Our experimental study duly confirms this assumption.

TABLE 1
Algebraic operations between types.

Operation	$\lvert\{n \mid (n_1, n) \in Q(G)\}\rvert$	$\lvert\{n \mid (n, n_2) \in Q(G)\}\rvert$	$\alpha(Q)$
$=$	Bounded	Bounded	0 or 1
$<$	Bounded	Not bounded	1
$>$	Not bounded	Bounded	1
$\diamond$	Not bounded	Not bounded	1
$\times$	Not bounded	Not bounded	2

5.2.2 Estimating the selectivity value

We propose a solution for estimating the selectivity value $\alpha(Q)$ of a given query Q, for all graphs satisfying a given schema $\mathcal{S}$. This basically means to compute a function that associates to Q a value $\hat{\alpha}(Q) \in \{0, 1, 2\}$. This value can be made more precise as follows: for a pair of node types A and B, $\hat{\alpha}_{A,B}(Q)$ is the estimated selectivity of Q restricted to pairs (x, y) where x is of type A and y of type B. Then, the overall estimated selectivity value of Q is $\hat{\alpha}(Q) = \max_{A,B}(\hat{\alpha}_{A,B}(Q))$.

To compute these values, we define an algebra based on what we call *selectivity classes*. First, for each node type A within the input schema $\mathcal{S}$, we denote $\mathsf{Type}(A) = N$ if A grows with the graph size and $\mathsf{Type}(A) = 1$ if it does not. In the graph schema, $\mathsf{Type}(A) = 1$ if $\mathcal{T}(A)$ is a fixed value and $\mathsf{Type}(A) = N$ if $\mathcal{T}(A)$ is a proportional value.

For each query Q and each pair of node types A and B, the *selectivity class of* Q *for* A, B, denoted $sel_{A,B}(Q)$ is a *triple* (t_A, o, t_B) such that $t_A = \mathsf{Type}(A)$, $t_B = \mathsf{Type}(B)$ and $o \in \{=, <, >, \diamond, \times\}$ is an operation between types.

We summarize these algebraic operations in Table 1, which should be read a follows: an operation from the first column denotes that for every graph G satisfying a schema $\mathcal{S}$, for every pair of nodes $(n_1, n_2) \in Q(G)$, it is the case that $|\{n \mid (n_1, n) \in Q(G)\}|$ and $|\{n \mid (n, n_2) \in Q(G)\}|$, resp., are or are not bounded (by some constants), as indicated in the second and third columns, resp. The last column $\alpha(Q)$ is particularly useful to distinguish between the last two operations $\diamond$ and $\times$.

We next intuitively explain the above operations and we illustrate them via examples:

• $=$ is the simplest operation and occurs either (i) between constant types e.g., (`country`, `language`) as illustrated for constant queries, or (ii) for some linear queries such as the query defined by the empty regular expression ε that returns precisely as many results as the nodes in the graph.

• $<$ characterizes queries where either (i) the out-degree distribution is Zipfian, such as (`language`, `user`) as illustrated for linear queries, or (ii) the source node type A has $\mathsf{Type}(A) = 1$ and the target node type B has $\mathsf{Type}(B) = N$. Then, the definition of $>$ is symmetric to $<$.

• $\times$ corresponds to queries performing a Cartesian product between two node sets (both growing with the graph), for example the transitive closure of the `knows` predicate that we used above to illustrate the quadratic queries. Intuitively, the $\times$ is the result of a $>$ followed by a $<$.

+	=	<	>	◇	×
=	=	<	>	◇	×
<	<	<	◇	◇	×
>	>	◇	>	◇	×
◇	◇	◇	◇	◇	×
×	×	×	×	×	×

(a) Disjunction.

·	=	<	>	◇	×
=	=	<	>	◇	×
<	<	<	×	×	×
>	>	◇	>	◇	×
◇	◇	◇	×	×	×
×	×	×	×	×	×

(b) Conjunction.

Fig. 5. Algebra for selectivity classes. To read in a (column, row) order.

• ◇ is the trickier operation and corresponds for instance to pairs of users that are known by someone in common. Most users are not linked in this way, but two pairs of hub users are. Thus, although there are numerous paths between two hubs, their number remains relatively small, in particular it grows linearly with the graph. Intuitively, the ◇ is the result of a $<$ followed by a $>$.

As defined in Section 3.3, the gMark workloads consist of queries expressed as UCRPQ's. Hence, we need to compute the selectivity values for regular path queries, which involve regular expressions. First, for a query Q defined by the regular expression ε (the empty word), for each type A, we have that $sel_{A,A}(Q) = (\mathsf{Type}(A), =, \mathsf{Type}(A))$. When Q is defined by a single edge label $a \in \Sigma^+$, we obtain $sel_{A,B}(Q)$ directly from the distribution of the a-labeled edges from A to B, as defined above and obtained from the schema.

Example 5.1 Consider the schema given in Example 3.3. First, we assign selectivity classes to the types of the schema, thus $\mathsf{Type}(T_1) = \mathsf{Type}(T_2) = N$ whereas $\mathsf{Type}(T_3) = 1$. From the schema, we compute the following values:

- $sel_{T_1,T_1}(a) = (N, <, N)$ and $sel_{T_1,T_1}(a^-) = (N, >, N)$ (because of the Zipfian out-distribution that moreover implies a Zipfian in-distribution for the inverse, and $\mathsf{Type}(T_1) = N$),
- $sel_{T_1,T_2}(b) = (N, =, N)$ and $sel_{T_2,T_1}(b^-) = (N, =, N)$ (because of non-Zipfian in- and out-distributions, and moreover, both $\mathsf{Type}(T_1) = \mathsf{Type}(T_2) = N$),
- $sel_{T_2,T_2}(b) = (N, =, N)$ and $sel_{T_2,T_2}(b^-) = (N, =, N)$ (same reasoning as for the previous bullet),
- $sel_{T_2,T_3}(b) = (N, >, 1)$ and $sel_{T_3,T_2}(b^-) = (1, <, N)$ (because of non-Zipfian in- and out-distributions, and moreover, $\mathsf{Type}(T_2) = N$ and $\mathsf{Type}(T_3) = 1$). ▫

Let a query Q be defined by the regular expression $p_1 + p_2$ where p_1 and p_2 are two regular expressions that define queries Q_1 and Q_2, respectively. For every pair of node types A, B, such that $sel_{A,B}(Q_1) = (t_A, o_1, t_B)$ and $sel_{A,B}(Q_2) = (t_A, o_2, t_B)$ then $sel_{A,B}(Q) = (t_A, o_1 + o_2, t_B)$ where $o_1 + o_2$ is defined by the table in Figure 5(a). Quite similarly, for Q defined by $p_1 \cdot p_2$, we have $sel_{A,B}(Q) = \Sigma_{C \in \Theta} sel_{A,C}(Q_1) \cdot sel_{C,B}(Q_2)$ where $sel_{A,C}(Q_1) \cdot sel_{C,B}(Q_2) = (t_A, o_1 \cdot o_2, t_B)$ for $sel_{A,C}(Q_1) = (t_A, o_1, t_C)$ and $sel_{C,B}(Q_2) = (t_C, o_2, t_B)$, and where $o_1 \cdot o_2$ is defined by the table in Figure 5(b). Finally, if Q is defined by the regular expression p^* where p defines a query Q', we assign a selectivity class to Q if and only if the input and output types of Q' are the same, in which case $sel_{A,A}(Q) = sel_{A,A}(Q') \cdot sel_{A,A}(Q')$.

As a remark, if either t_A or t_B is 1, the operator solely relies on the other one. Hence, the triples $(1, \times, 1)$ and $(1, \diamond, 1)$ are not permitted, which makes $(1, =, 1)$, $(1, <, N)$ and $(N, >, 1)$ the only permitted triples that contain a 1. However, in the computation of the algebraic expression, we could still obtain triples $(1, \times, 1)$ and $(1, \diamond, 1)$, that we should replace with $(1, =, 1)$ if the case occurs.

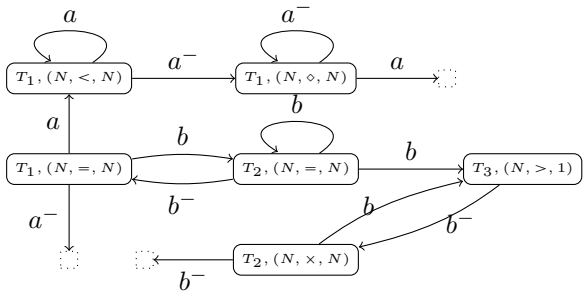

Fig. 6. A snippet of the schema graph for our running example.

Finally, the *estimated selectivity value* of a query is obtained directly from its class. If the obtained selectivity class of a query Q is $(1, =, 1)$, then $\hat{\alpha}(Q) = 0$, if we obtain $(N, \times, N)$, then $\hat{\alpha}(Q) = 2$, and $\hat{\alpha}(Q) = 1$ for all other cases.

5.2.3 Data structures

The data structures needed to estimate the selectivity are: (a) the schema graph $\mathbb{G}_{\mathcal{S}}$, (b) the distance matrix $\mathcal{D}$, and (c) the selectivity graph $\mathbb{G}_{sel}$.

(a) Schema graph $\mathbb{G}_{\mathcal{S}}$. We derive the schema graph from the schema $\mathcal{S}$. Each node in $\mathbb{G}_{\mathcal{S}}$ is a pair given by a node type of the schema and a selectivity triple (cf. Section 5.2.2) associated to that type. More formally, the set of nodes of the schema graph $\mathbb{G}_{\mathcal{S}}$, denoted $SelType(\mathcal{S})$, consists of tuples $(T, (t_1, o, \mathsf{Type}(T)))$, where (i) T is a node type from Θ and (ii) $(t_1, o, \mathsf{Type}(T))$ is a selectivity triple in the set of all possible selectivity triples $\{(1, =, 1), (1, <, N), \ldots\}$. The edges of $\mathbb{G}_{\mathcal{S}}$ are labeled with symbols in Σ^+.

The goal of the schema graph $\mathbb{G}_{\mathcal{S}}$ is to indicate how a path ending with a type T of selectivity triple $(t_1, o, \mathsf{Type}(T))$ changes when it is extended with an edge in Σ^+. Formally, given a node $(T, (t_1, o, \mathsf{Type}(T)))$ and a label (or label inverse) $a \in \Sigma^+$ such that the schema allows an a-labeled edge between T and a node type T', if according to our algebra we have $(t_1, o, \mathsf{Type}(T)) \cdot sel_{T,T'}(a) = (t_1, o', \mathsf{Type}(T'))$, then in the schema graph $\mathbb{G}_{\mathcal{S}}$ there is an edge $((T, (t_1, o, \mathsf{Type}(T))), a, (T', (t_1, o', \mathsf{Type}(T'))))$.

Example 5.2 Recall Examples 3.3 and 5.1. We illustrate a snippet of the corresponding schema graph in Figure 6. For instance, the node $(T_1, (N, =, N))$ is due to the fact that $\mathsf{Type}(T_1) = N$ and recall from Section 5.2.2 that for a given type A we have $sel_{A,A}(\varepsilon) = (\mathsf{Type}(A), =, \mathsf{Type}(A))$. Moreover, our schema allows an a-labeled edge between two nodes of type T_1, following a Zipfian out-degree distribution, hence its selectivity triple is $(N, <, N)$, which explains the node $(T_1, (N, <, N))$ in Figure 6. Additionally, there is an a-labeled edge in our schema graph between the nodes $(T_1, (N, =, N))$ and $(T_1, (N, <, N))$ because in our algebra $(N, =, N) \cdot (N, <, N) = (N, <, N)$. ▫

(b) Distance matrix $\mathcal{D}$. The distance matrix $\mathcal{D}$ establishes for each pair $n, n' \in SelType(\mathcal{S})$ the length $\mathcal{D}(n, n')$ of the shortest path between n and n' in $\mathcal{D}(n, n')$ in $\mathbb{G}_{\mathcal{S}}$.

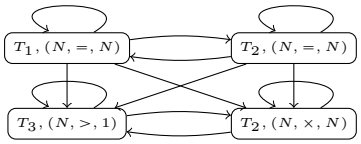

Fig. 7. An excerpt of the selectivity graph derived from the schema graph in Figure 6.

(c) Selectivity graph $\mathbb{G}_{sel}$. The selectivity graph $\mathbb{G}_{sel}$ is an unlabeled directed graph whose nodes are $SelType(\mathcal{S})$. An edge exists between two nodes n and n' if there exists a path between n and n' of length in $[l_{\min}, l_{\max}]$ (recall that this interval is specified as part of the query size in the workload configuration $\mathbb{Q}$).

Example 5.3 We illustrate the selectivity graph for our running example in Figure 7. As an example, there exists a path between $T_1, (N, =, N)$ and $T_2, (N, \times, N)$ (for instance, the path $b \cdot b \cdot b^-$), whose length is less than $l_{\max} = 4$. However, there does not exist such a path between $T_2, (N, \times, N)$ and $T_1, (N, =, N)$, thus there is no edge between those nodes. ▫

5.2.4 Accommodating selectivity in query generation

We show how the instantiation of the placeholders (line 4 in Algorithm 2) should be implemented to support selectivity estimation for binary queries. Assume that we have already generated the query skeleton and picked the two projection variables (lines 2-3).

We first build a function T that associates to each placeholder P of the query skeleton a *selectivity type* $(T_1, (\mathsf{Type}(T_1), o, \mathsf{Type}(T_2)), T_2)$, where T_1 and T_2 are schema types, and o is a selectivity operator. This function exhibits two properties. First, the input and output type of each selectivity operator should be consistent with the types of T_1, T_2 and secondly, selectivity values for each type should be selected in a way to guarantee that the global selectivity class for the query is the expected one. To achieve this, we randomly choose a path on the selectivity graph $\mathbb{G}_{sel}$. This basically means that we want to find a path between a node with selectivity triple $(?, =, ?)$ (where by "?" we denote any type) to a node with (T_1, o, T_2) yielding one of the desired selectivities. The length of this path should be consistent with the configuration. Some of the placeholders contain Kleene stars, with a probability dictated by the configuration. Such conjuncts inherit the input and output types of their neighbor conjuncts, with the selectivity operator '='.

Example 5.4 In our running example, assuming that we look for a linear query with 3 conjuncts, we can instantiate the function T as follows:

$$T(P_1) = (T_1, (N, =, N), T_1), \quad T(P_2) = (T_1, (N, >, N), T_2),$$
$$T(P_3) = (T_2, (N, =, N), T_2).$$

We can then compute the selectivity of the concatenation, which is $(N, =, N) \cdot (N, >, N) \cdot (N, =, N) = (N, >, N)$, which corresponds to a linear query. ▫

Note that drawing uniformly at random paths of a certain length in $\mathbb{G}_{sel}$ can be done efficiently with a two-step algorithm: first, each node n is associated with a function $nb\_path(n, i)$ that gives the number of paths of length i that can be generated starting from n. For instance, for a quadratic query, a node $n_1 : (A, (N, \times, N))$ has $nb\_path(n_1, 0) = 1$ whereas a node $n_2 : (B, (1 = 1))$ has $nb\_path(n_2, 0) = 0$. Other values are obtained by a saturation algorithm: to generate a path of length l, the algorithm picks a starting node with a random draw weighted by $nb\_path(n, l)$, and then picks the label of an outgoing edge to a node n' with a random draw weighted by $nb\_path(n', l-1)$, etc. until all the nodes are saturated.

Next, for each placeholder we build a path skeleton that satisfies the constraints concerning the number of disjuncts and the length of the paths in the query size, as illustrated by the following example.

Example 5.5 To continue our running example, with a number of disjuncts in the range $[d_{\min}, d_{\max}] = [3, 5]$ and path length in the range $[l_{\min}, l_{\max}] = [2, 4]$, we may have a path skeleton as follows: $P_1 = X_1 \cdot X_2 + X_3 \cdot X_4 \cdot X_5 + X_6 \cdot X_7$ (i.e., three disjuncts having path length from 2 to 3). ▫

Once a path skeleton is computed, we need to find actual edge labels for each path, by randomly choosing paths in $\mathbb{G}_{sel}$. For the sake of conciseness, we omit the details and we illustrate it on the following example.

Example 5.6 Let us consider $P_1 = X_1 \cdot X_2 + X_3 \cdot X_4 \cdot X_5 + X_6 \cdot X_7$, with $T(P_1) = (T_1, (N, =, N), T_1)$. The query workload generation algorithm may obtain a path instantiation of this kind: $X_1 = b$ and $X_2 = b^-$ as the path $b \cdot b^-$ can go from T_1 to T_2 and from T_2 to T_1 via a concatenation of two paths with edge label b. ▫

The total size of the data structures described in Section 5.2.3 is quadratic in the size of the graph schema (due to the distance matrix) and the running time of the selectivity estimation algorithm inherits the quadratic behavior. However, we observed that the query workload generation is very efficient, as we point out in more details at the end of Section 6.2 for our default schemas and for the three existing schemas that we encoded in gMark. Moreover, it may not be possible that all constraints are satisfied at the same time and testing satisfiability is intractable (cf. Section 3.4). Our algorithm always outputs a result to the user. More precisely, when instantiating the placeholders, it may be the case that we cannot fill the path skeleton with paths of the precise lengths and meeting the required selectivities, hence we choose to relax the path length in order to ensure accurate selectivity estimation and efficiency (instead of backtracking and drawing new skeletons).

6 Empirical evaluation of gMark

In this section, we empirically evaluate gMark w.r.t. two important aspects: the capability to encode the application domain of existing benchmarks (Section 6.1), and its quality in terms of both accuracy of the estimated selectivities and scalability of the generator (Section 6.2).

Environment. All experiments reported in this section and in Section 7 were run on an Intel Core i7 920, with 6GB RAM, and running Ubuntu 14.04 64bit.

6.1 Coverage of practical graph scenarios

In our experiments, we relied on four use cases: the default gMark use case, and three gMark encodings of the schemas of existing state-of-the-art benchmarks [5]–[7] that were possible since gMark can be easily tuned to fit an arbitrary set of predicates, node types, and schema constraints.

• `Bib` is our default scenario, describing the bibliographical database introduced in Section 3.1 as our motivating example. It represents a baseline, illustrating the main gMark features, in particular all types of degree distributions.

• `LSN` is our gMark encoding of the fixed schema provided with the LDBC Social Network Benchmark [7], which simulates user activity in a social network.

• `SP` is our gMark encoding of the fixed DBLP-based [28] schema provided with SP2Bench [6].

• `WD` is our gMark encoding of the default schema provided with Waterloo SPARQL Diversity Test Suite (WatDiv) [5], which differs from LDBC and SP2Bench, and is similar to gMark in the sense that it also supports user-defined schemas via a data set description language. The default schema that we encoded is about users and products.

The difference between gMark and the aforementioned benchmarks resides in the kind of expressible schema constraints that are incomparable. Nonetheless, we have been able to encode their key characteristics, which include for instance node types, edge labels, and associations between entities. By opposite, we could not encode other characteristics that are typical of those benchmarks, such as subtyping and hardcoded correlations that we do not support in gMark. We detail in Appendix A in our technical report [29] the expressiveness differences between gMark and existing benchmarks, and our encoding choices of their schemas.

Discussion on the query loads. Because of these differences, the gMark query loads are different from the fixed query loads in the other benchmarks. Nevertheless, although gMark does not generate exactly the same queries of the other benchmarks, it can easily be tuned to generate queries with similar characteristics, i.e., queries of comparable query shape and size and featuring the same selectivity of the original queries of those benchmarks. To concretely point out that this is the case, we report in Figure 8 an experiment in which we considered the execution times of three queries (one per selectivity class) from the original SP2Bench query load and three comparable queries of the same shape, size and selectivity generated with gMark using `SP`. We observe that the queries generated and executed with gMark show the same asymptotic runtime behavior of the original SP2Bench queries executed under SP2Bench. Precisely, we can notice in Figure 8 that the obtained simulated constant (linear and quadratic, respectively) query in gMark falls in the same selectivity class of the original constant (linear and quadratic, respectively) query in SP2Bench. Since the goal of gMark is not to simulate the exact query loads of other existing benchmarks, in the rest of the experimental study we only rely on the gMark encodings of their schemas `LSN`, `SP`, `WD` (as well as our user-defined schema `Bib`) to generate more diverse query loads, going beyond the existing ones w.r.t. the fine-grained control of different characteristics such as size, selectivity, and recursion.

TABLE 2
$\alpha(Q)$ averaged across constant, linear, and quadratic queries (with standard deviation), with varying graph sizes, data, and query diversity.

	Constant	*Linear*	*Quadratic*
`LSN`-*Len*	0.200±0.417	1.189±0.261	2.032±0.059
`LSN`-*Dis*	0.182±0.364	1.325±0.318	2.046±0.074
`LSN`-*Con*	0.190±0.391	1.244±0.326	2.017±0.032
`LSN`-*Rec*	0.196±0.409	1.090±0.492	1.564±0.889
`Bib`-*Len*	0.003±0.010	0.921±0.122	1.405±0.337
`Bib`-*Dis*	0.000±0.000	0.995±0.012	1.607±0.261
`Bib`-*Con*	0.023±0.029	0.986±0.112	1.409±0.296
`Bib`-*Rec*	0.100±0.316	0.982±0.073	1.493±0.335
`WD`-*Len*	0.016±0.044	1.427±0.392	2.004±0.022
`WD`-*Dis*	0.009±0.022	1.412±0.380	1.999±0.014
`WD`-*Con*	-0.010±0.026	1.540±0.495	1.750±0.708
`WD`-*Rec*	0.587±0.830	-	1.976±0.012
`SP`	0.074±0.130	1.064±0.034	2.034±0.295

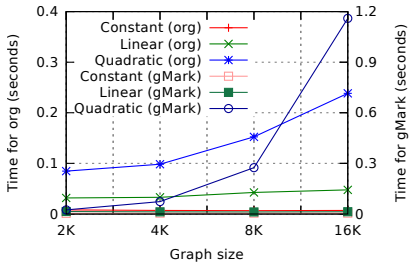

Fig. 8. Evaluation times of a constant, a linear and a quadratic query in the original SP2Bench query load (org) and gMark, respectively.

6.2 Quality and scalability

Selectivity estimation quality study. Our first set of experiments focuses on understanding the quality of the selectivity estimation performed using the algebra presented in Section 5.2. We used four types of workloads in which we stress-test several diverse queries with incrementally varied query size: *Len* generates queries with varying path lengths, no disjuncts, no conjuncts, and no recursion; *Dis* generates queries with disjuncts, no conjuncts and no recursion; *Con* generates queries with conjuncts and disjuncts and no recursion; *Rec* generates queries with recursion (Kleene-stars). Each of the above tests leads to a workload of 30 queries, out of which 10 are constant, 10 are linear, and 10 are quadratic. We repeated these experiments for each of the `Bib`, `WD`, and `LSN` use cases (cf. Section 6.1). We executed each query Q on graph instances of sizes between 2K and 32K, and we counted the number of results returned on each instance. To compute the α-value in the formula $|Q(G)| = \beta|G|^{\alpha}$ (cf. Section 5.2), we computed a simple linear regression between $\log|G|$ and $\log|Q(G)|$. We averaged the α-values obtained across all queries belonging to the same selectivity class. We report these results, together with the standard deviation, in Table 2. As expected, we observe that, for all combinations of use cases and diverse query workloads, the estimated values of α are ≈ 0 for constant queries, ≈ 1 for linear queries, and ≈ 2 for quadratic queries. The missing value in Table 2 corresponds to linear recursive queries for `WD`. In the presence of recursion, we actually observed numerous failures on the majority of the studied systems (cf. Section 7).

Finally, the last row (`SP`) of Table 2 shows the estimated α-values for SP2Bench [6] on a set of queries following our gMark encoding of the original set of SP2Bench queries. In the remainder of our study in Section 7, we disregard `SP`

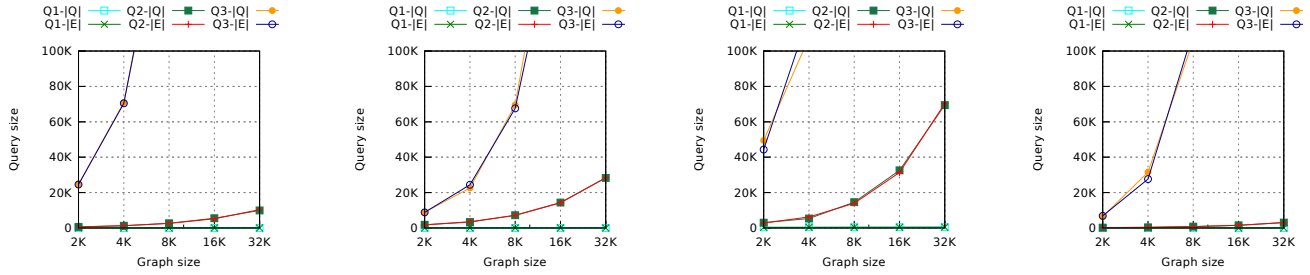

(a) Est. selectivities: Bib-*Len.* (b) Est. selectivities: Bib-*Con.* (c) Est. selectivities: Bib-*Dis.* (d) Est. selectivities: Bib-*Rec.*

Fig. 9. Summary of selectivity estimation quality results for the bibliographical use case Bib.

TABLE 3
Graph generation time for varying graph sizes (# nodes) and schemas.

	100K	1M	10M	100M
Bib	0m0.057s	0m0.638s	0m8.344s	1m28.725s
LSN	0m0.225s	0m1.451s	0m23.018s	3m11.318s
WD	0m2.163s	0m25.032s	4m10.988s	113m31.078s
SP	0m0.638s	0m7.048s	1m28.831s	15m23.542s

since this use case does not bring more insights than the query sets generated for the other use cases.

To further illustrate the precision of our estimated values, we report in Figure 9 the estimated selectivities ($|E|$) along with the theoretical selectivities ($|Q|$) for constant (Q_1), linear (Q_2), and quadratic (Q_3) queries on the Bib use case. We observe that for the classes of queries of increasing expressiveness the number of results is generally higher for quadratic queries, while it is linearly and constantly varying for the other queries, as expected. We also observe that the two curves representing the estimated selectivities and the theoretical ones closely overlap in all the cases. Finally, notice that the above experiments are considering chain queries only. The results on other shapes and/or use cases are similar and omitted for conciseness.

We conclude from this study that the schema-driven gMark selectivity estimation framework generates consistently high quality estimates, across all selectivity classes, across a broad spectrum of diverse data sets and queries.

Scalability study. Our second set of experiments is devoted to measuring the time taken by the graph generator of gMark, while varying the size of the data and the size of the query workload. To gauge the robustness of our graph instance generator w.r.t. data and query diversity, in these experiments we employed all four use cases: Bib, WD, LSN, and SP. We report the results in Table 3. We observe that the generator scales quite well for all use cases. It is quite efficient for big graph sizes in all cases except WD. This is due to the quite complex nature of its schema, which induces much denser graph instances compared to the other use cases. For example, WD instances have two orders of magnitude higher number of edges than Bib instances having the same number of nodes. This is not a limitation of gMark or of the WD scenario, but rather a specific feature of very dense graphs.

We conclude by observing that gMark can efficiently generate both small and large graph instances, on a diversity of practical scenarios. We note that we also conducted a scalability study of query generation, which showed that gMark easily generates workloads of a thousand queries for Bib, LSN, and SP in around one second and for the richer WD scenario in around 10 seconds. Query translation of a thousand queries into all four supported syntaxes for each of the four scenarios took a mere tenth of a second. This study shows that gMark workload generation is very efficient and scalable for large-scale complex workloads.

7 Evaluation of Query Engines

We next turn to an empirical evaluation of a representative selection of currently available graph query processing engines using gMark. Our goal here is both to demonstrate the new capabilities in benchmarking introduced by gMark, and to pinpoint limitations and areas for further improvement in current graph query processing solutions.

7.1 Design of experiments

Systems. The database systems (and their supported query languages) that we consider in our study are[3]:

- **G**: a native graph database (openCypher [30])
- **S**: a popular SPARQL query engine (SPARQL 1.1 [31])
- **P**: PostgreSQL v9.3.9 (SQL:1999 recursive views[4] [32])
- **D**: a modern Datalog engine (Datalog [33])

For the sake of fairness, we used default configurations for each system i.e., without special purpose optimizations.

Query languages. Recall that the queries generated by gMark are UCRPQ's. We provide in Appendix B in our technical report [29] a translation of an example UCRPQ into each of the above concrete syntaxes. We note that not all systems support the full expressive power of UCRPQ's. In particular, arbitrary UCRPQ's can be expressed in SPARQL, SQL, and Datalog, while openCypher supports only those UCRPQ's having no occurrences of inverse or concatenation under Kleene star. In our results regarding recursive queries, some of the generated benchmark queries do indeed exhibit inverse and/or concatenation in a recursive conjunct. In these cases, the corresponding openCypher query has only the non-inverse symbol and/or the first symbol in a concatenation of symbols, respectively. Furthermore, while all other languages adopt the classical homomorphic semantics for conjunctive queries [33], openCypher adopts an isomorphic semantics. For these two reasons, openCypher queries often

3. For obvious reasons, we obfuscate the names of the three commercial systems employed in our study.

4. We use the standard translation of UCRPQ's into recursive views, implemented using linear recursion [26].

have answer sets that differ from that of their counterparts in the other languages, which should be kept in mind while evaluating experimental results pertaining to system **G**.

Measurements. We generate and execute query workloads on a variety of graph configurations. We execute and measure the runtime of each query six times. The first one is a "cold" run that we exclude from the computation of the average; from the remaining five "warm" runs we drop the fastest and slowest and then report the average of the remaining three execution times. Between the execution of each query, we close and reopen the database to clear any caching effects. We consider the following parameters in generating workloads: selectivity classes of the tested queries (constant, linear, quadratic) and size (amounting to 30 queries for each workload). We consider the following parameters in generating graphs: size (from 2K to 16K nodes) and use case i.e, `Bib`, `LSN`, and `WD` (cf. Section 6.1).

We make three remarks here. (i) Despite being small, the considered graph sizes were already sufficient to illustrate interesting behavior and distinctions between the studied systems. Indeed, as we discuss in Section 7.2, already on these graphs we observed that many simple queries fail on a majority of the systems. Even for those queries that succeed, the evaluation times are often very high e.g., hundreds or even thousands of seconds already on instances of these sizes. (ii) To ensure a fair comparison of all systems and to avoid measuring the time to print the query results, we added to all queries the aggregate $count(distinct(\overline{?v}))$, where $\overline{?v}$ is the (binary) vector of output variables. We recall that *distinct* is also necessary for our analysis since the algebra relies on the elimination of duplicates (cf. Section 5). (iii) In this study, all queries are chains, as this is the basic query shape from which the others are constructed and hence sufficient for illustrating the relative performance of current graph query processing engines. Finally, even though in the presentation of the results, we focus on the default use case `Bib`, we observed comparable trends for the other use cases, that we omit for the sake of conciseness.

7.2 Results of experiments

Non-recursive queries. In our first experiment, we focus on the non-recursive queries i.e., query workloads *Len*, *Dis*, *Con* (cf. Section 6.2). We summarize the results in Figure 10. The goal of this study is to observe how the different systems react to these varied workloads. Figure 10(a) shows the query execution times averaged across the 10 constant queries of each workload. Out of the 10 averages obtained by the 5 warm runs of each query, we computed once again the overall average, by discarding two out of the 10 averages that have the farthest standard deviation with respect to this overall average. This allows to capture the cases in which some of the systems fail or give outlier results. In Figure 10(a), we observe that **P** reacts better than **S**, **G**, and **D** to query workload diversity, by exhibiting lower query evaluation times on all instance sizes. This behavior is confirmed in the case of linear queries, as shown in Figure 10(b), for the *Con* query workload on all the sizes and for the *Len* and *Dis* query workloads for smaller sizes only i.e., 2K and 4K. For larger instance sizes i.e., 8K and 16K, the behavior is reverted in favor of **S** for linear queries. Then, as shown in Figure 10(c) for quadratic queries, **S** continues on this trend by beating **P**, **G**, and **D**. We also observe from all query execution times reported in Figure 10 that the times taken by constant and linear are of the same order of magnitude, whereas quadratic queries, as expected, typically exhibit an order of magnitude slowdown. There is only one system (**D**) for which the differences of the behavior along the sets of linear and quadratic queries are blurred.

We conclude from this study that gMark allows us to generate interesting queries and diverse query workloads that already on small graph instances stress-test state-of-the-art systems, and highlight particular strengths and weaknesses in processing graph queries. As a general observation, we can further conclude that the straightforward standard implementation of UCRPQ's in PostgreSQL typically shows superior performance across a broad class of queries (i.e. constant and linear) to that of existing dedicated systems.

TABLE 4
Execution time (sec.) for recursive queries.

Syst.	*Query 1: Graph Size*				*Query 2: Graph Size*			
	2K	4K	8K	16K	2K	4K	8K	16K
P	3400	72113	-	-	-	-	-	-
G	-	-	-	-	-	-	-	-
S	6621	-	-	-	-	-	-	-
D	450	455	552	725	607	704	1295	2095

Recursive queries. Our second experiment is devoted to recursive queries, generated by query workloads containing Kleene stars. Unfortunately, all tested systems either failed on the majority of these queries or had to be manually terminated after unexpectedly long running times. For these reasons, it is difficult to draw a clear conclusion on recursive queries. Therefore, we performed a small case analysis: we considered two recursive queries of constant and quadratic selectivity, respectively, for which we could collect results for at least one of the four systems. We report the results for both queries in Table 4. The first query has constant selectivity. **P** was quite slow at evaluating it on small instances and starts failing on graphs of 8K nodes. **S** was able to answer this query only on the smallest graph size (2K). **G** failed in all cases and always returned empty results (due to its different query semantics, as discussed in Section 7.1). The only system for which we could measure the evaluation time for all sizes was **D**, which also turned to be the most efficient one. The second query has quadratic selectivity and only **D** was able to evaluate it. We conclude from this study that only **D** is currently able to deal with recursive queries.

8 Concluding Remarks

We presented gMark, the first generator that satisfies the key criteria of being domain-independent, extensible, schema-driven, and highly configurable also in terms of the expected query selectivity of a given workload. The latter is a novel contribution on its own and is applicable to other independent benchmarks and problems. For instance, we could envision the query workload generation in gMark applied to real graph data sets on top of which a schema extraction tool has been run beforehand. Furthermore, gMark is the first benchmark to generate workloads exhibiting recursive path queries, which are central to graph querying. Our in-depth

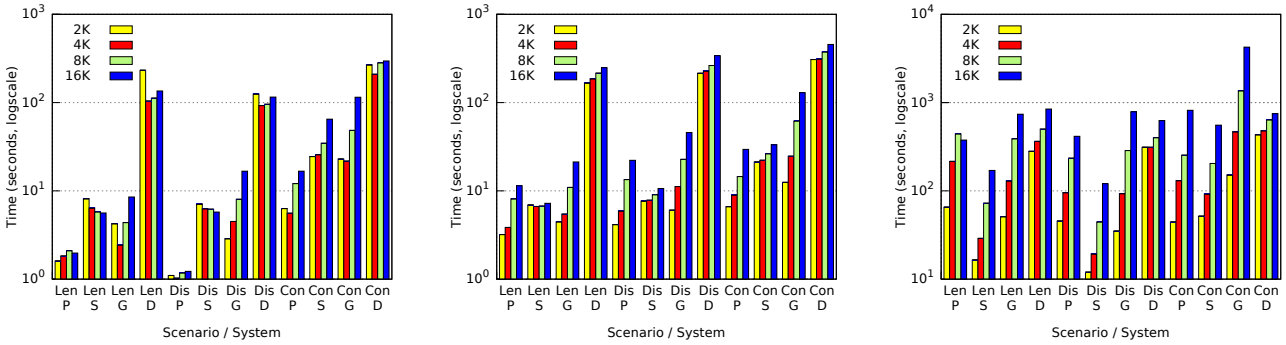

(a) *Constant* queries. (b) *Linear* queries. (c) *Quadratic* queries.

Fig. 10. Summary of query execution times for diverse query workloads (*Len*, *Dis*, *Con*) and various graph sizes under PostgreSQL (**P**) and three commercial systems: a SPARQL query engine (**S**), a native graph database (**G**), and a modern Datalog engine (**D**).

empirical study demonstrated both the quality and practicality of gMark. Moreover, our experiments highlighted important limitations in the query processing capabilities of current state-of-the-art graph processing engines, already on small graph instances and on both recursive and non-recursive queries.

We plan to align our work on gMark with international benchmarking bodies such as LDBC [7]. Looking ahead to the future work, there are many directions for further investigation e.g., extending the selectivity estimation to n-ary queries. We also aim to evangelize for the use of gMark by researchers in the graph data management community; a first step in this direction is our VLDB demo [34].

REFERENCES

[1] http://neo4j.com .
[2] http://www.sparsity-technologies.com .
[3] M. Aref, B. ten Cate, T. J. Green, B. Kimelfeld, D. Olteanu, E. Pasalic, T. L. Veldhuizen, and G. Washburn, "Design and implementation of the LogicBlox system," in *SIGMOD*, 2015, pp. 1371–1382.
[4] http://virtuoso.openlinksw.com .
[5] G. Aluç, O. Hartig, M. T. Özsu, and K. Daudjee, "Diversified stress testing of RDF data management systems," in *ISWC*, 2014, pp. 197–212.
[6] M. Schmidt, T. Hornung, G. Lausen, and C. Pinkel, "SP2Bench: A SPARQL performance benchmark," in *ICDE*, 2009, pp. 222–233.
[7] O. Erling, A. Averbuch, J. Larriba-Pey, H. Chafi, A. Gubichev, A. Prat, M.-D. Pham, and P. Boncz, "The LDBC social network benchmark: Interactive workload," in *SIGMOD*, 2015, pp. 619–630.
[8] D. Barbosa, I. Manolescu, and J. X. Yu, "XML benchmarks," in *Encyclopedia of Database Systems*, L. Liu and M. T. Özsu, Eds. Springer, 2009, pp. 3576–3579.
[9] J. Gray, Ed., *The Benchmark Handbook for Database and Transaction Systems (2nd Edition)*. Morgan Kaufmann, 1993.
[10] Y. Guo, Z. Pan, and J. Heflin, "LUBM: A benchmark for OWL knowledge base systems," *J. Web Sem.*, vol. 3, no. 2-3, pp. 158–182, 2005.
[11] C. Bizer and A. Schultz, "The Berlin SPARQL benchmark," *Int. J. Semantic Web Inf. Syst.*, vol. 5, no. 2, pp. 1–24, 2009.
[12] D. Blum and S. Cohen, "Grr: Generating random RDF," in *ESWC*, 2011, pp. 16–30.
[13] S. Duan, A. Kementsietsidis, K. Srinivas, and O. Udrea, "Apples and oranges: a comparison of RDF benchmarks and real RDF datasets," in *SIGMOD*, 2011, pp. 145–156.
[14] S. Qiao and Z. M. Özsoyoglu, "RBench: Application-specific RDF benchmarking," in *SIGMOD*, 2015, pp. 1825–1838.
[15] J. W. Zhang and Y. C. Tay, "GSCALER: Synthetically scaling a given graph," in *EDBT*, 2016, pp. 53–64.
[16] http://snap.stanford.edu/ .
[17] http://konect.uni-koblenz.de .
[18] S. Staworko, I. Boneva, J. E. L. Gayo, S. Hym, E. G. Prud'hommeaux, and H. R. Solbrig, "Complexity and expressiveness of ShEx for RDF," in *ICDT*, 2015, pp. 195–211.
[19] A. Arasu, R. Kaushik, and J. Li, "Data generation using declarative constraints," in *SIGMOD*, 2011, pp. 685–696.
[20] N. Bruno, S. Chaudhuri, and D. Thomas, "Generating queries with cardinality constraints for DBMS testing," *IEEE Trans. Knowl. Data Eng.*, vol. 18, no. 12, pp. 1721–1725, 2006.
[21] A. Gubichev and P. Boncz, "Parameter curation for benchmark queries," in *TPCTC*, 2014, pp. 113–129.
[22] E. Lo, C. Binnig, D. Kossmann, M. T. Özsu, and W. Hon, "A framework for testing DBMS features," *VLDB J.*, vol. 19, no. 2, pp. 203–230, 2010.
[23] E. Lo, N. Cheng, W. W. K. Lin, W. Hon, and B. Choi, "MyBenchmark: generating databases for query workloads," *VLDB J.*, vol. 23, no. 6, pp. 895–913, 2014.
[24] C. Mishra, N. Koudas, and C. Zuzarte, "Generating targeted queries for database testing," in *SIGMOD*, 2008, pp. 499–510.
[25] M. Poess and J. M. Stephens, "Generating thousand benchmark queries in seconds," in *VLDB*, 2004, pp. 1045–1053.
[26] P. Barceló, "Querying graph databases," in *PODS*, 2013, pp. 175–188.
[27] T. J. Schaefer, "The complexity of satisfiability problems," in *STOC*, 1978, pp. 216–226.
[28] http://dblp.uni-trier.de/db/ .
[29] G. Bagan, A. Bonifati, R. Ciucanu, G. Fletcher, A. Lemay, and N. Advokaat, "gMark: Schema-driven generation of graphs and queries," 2015. http://arxiv.org/abs/1511.08386.
[30] http://www.opencypher.org/ .
[31] http://www.w3.org/TR/sparql11-overview/ .
[32] http://www.postgresql.org .
[33] J. D. Ullman, *Principles of database and knowledge-base systems, Volumes I & II*. Computer Science Press, 1988, 1989.
[34] G. Bagan, A. Bonifati, R. Ciucanu, G. H. L. Fletcher, A. Lemay, and N. Advokaat, "Generating flexible workloads for graph databases," *PVLDB*, vol. 9, no. 13, 2016.

Guillaume Bagan (PhD, University of Caen, 2009) is a CNRS research engineer at the University of Lyon 1. His research focuses on query evaluation for relational and graph databases.

Angela Bonifati (PhD, Politecnico di Milano, 2002) is a professor at the University of Lyon 1, conducting research on graph databases, data integration for complex data formats, and query inference.

Radu Ciucanu (PhD, University of Lille 1 and Inria, 2015) is an associate professor at the University Blaise Pascal, Clermont-Ferrand, working mainly on graph databases and data integration.

George Fletcher (PhD, Indiana University Bloomington, 2007) is an associate professor at TU Eindhoven. He studies data-intensive systems, with a focus on the theory and engineering of query languages.

Aurélien Lemay (PhD, University of Lille 3, 2002) is an associate professor at the University of Lille 3 and Inria. His research focuses on language theory, database theory, and grammatical inference.

Nicky Advokaat obtained a MSc in Computer Science in 2015 at TU Eindhoven, with a thesis on graph database benchmarking.

APPENDIX

A Details on use cases

In this section, we provide further details on the gMark encoding of the LDBC [7], SP2Bench [6] and WatDiv [5] benchmarks (introduced in Section 6.1) and discuss the adequacy of gMark for modeling them as graph configurations.

LDBC. The LDBC Social Network Benchmark [7] has a fixed schema that simulates user activity in a social network. Its schema consists of 11 entities (such as *Persons*, *Tags*, *Forums*, *Messages*, *Likes*, *Organizations*, *Places*, etc.) and 20 relations (such as *knows*, *studyAt*, *workAt*, *likes*, *hasMember*, etc.). There are three main differences between LDBC and gMark: (i) LDBC relies on the aforementioned fixed schema, while gMark can be configured to fit an arbitrary schema, (ii) LDBC supports subtyping (e.g., *City*, *Country*, and *Continent* are all subtypes of the type *Place*) whereas this kind of feature is out of the scope of gMark, and (iii) LDBC encodes the associations between the different entities via standard UML notation (such as one-to-many or many-to-many associations) whereas gMark can finely tune these associations with in- and out-degree distributions,. Consequently, since the LDBC social network benchmark and gMark are incomparable, we had to make some choices when modeling LDBC in our framework. More precisely, regarding (i) and (ii) above, we have not specified in gMark the types that have subtypes but rather have specified directly the subtypes (e.g., we do not consider the type *Place* but we consider instead all of its subtypes *City*, *Country*, and *Continent*). More importantly, regarding (iii), we used our fine-grained support for encoding distributions between entities to replace generic associations with more natural distributions between them. For example, we encoded the relation *knows* between two persons with a Zipfian on both the in- and out- distributions (since in a real-world network there are indeed some authorities that know and are known by many others, while the majority of people know and are known by few others), and we encoded *workAt* with a Zipfian in-distribution and a uniform out-distribution (since intuitively a person works for a small number of companies, while there are few companies with a lot of employees and more companies with few employees). To sum up, typical real-world associations such as *knows* and *workAt*, encoded simply as many-to-many relationships in LDBC, can be modeled with finer grain by using the gMark bidirectional distribution control.

SP2Bench. The SP2Bench [6] benchmark has a fixed DBLP-based schema [28]. The only parameter that the user can control in the graph generator is the graph size (given either as number of tuples or number of years with publications), and the generator hardcodes a large number of constraints to make the generated graphs resemble as much as possible with the real DBLP. On the other hand, the user can declaratively specify in gMark much more parameters than only the size of the graph. Due to its generality, gMark cannot encode very specific characteristics such making the parameters μ/σ for the Gaussian distribution depend on the value of the publication year as SP2Bench does. Nonetheless, gMark easily captures the essential properties, such as the power-law distributions (found in the citation system or the distribution of papers among authors). Moreover, SP2Bench contains constraints like $x\%$ of the nodes of type A have an attribute B, that we encode in gMark as follows: we specify how many A's and B's we want in the graph, and the distributions between A and B. From the query workload point of view, SP2Bench was designed specifically to cover as much as possible from the SPARQL operators, whereas we are independent from the concrete query language. At the moment when SP2Bench was designed, the SPARQL specification did not include yet the property paths, whereas gMark covers such a feature. Moreover, the SP2Bench designers relied on the knowledge of the schema to be able to manually design a few meaningful queries with different characteristics, whereas gMark automatically generates a controlled diverse workload for an arbitrary schema.

WatDiv. The Waterloo SPARQL Diversity Test Suite (WatDiv) [5] differs from LDBC and is similar to ours in the sense that it also supports user-defined schemas via a so-called data set description language. However, the kind of schema constraints expressible in WatDiv and gMark differ. Indeed, WatDiv supports only local constraints for the predicates (that are however different from ours), while gMark supports global constraints on both types and predicates. More precisely, the local constraints in WatDiv specify for instance that 30% of the products have a content rating (we used a similar encoding as explained above for SP2Bench for such constraints), whereas gMark can encode that there is a Zipfian in- and out-distribution between product and rating. Additionally, gMark can specify global constraints also on the predicates (e.g., 40% of the edges should be labeled *friendOf*) whereas WatDiv does not. Another difference is that WatDiv relies solely on RDF while our model is general enough to support any graph encoding. However, since they rely on RDF, they also make use of subtyping similar to LDBC (e.g., a product can be an album, a movie, a concert, etc.). For the goal of our study, we encoded the standard WatDiv test data set schema (based on products, retailers, users, etc.) while using only the subtypes instead of general types and adapting their local constraints to our semantics.

B Examples of query translations

In this section, we give an example query in UCRPQ along with its translations into all four practical query languages supported by gMark, namely SPARQL, Datalog, openCypher, and SQL. The query that we discuss here contains one query rule, and consists of both non-recursive and recursive conjuncts, along with disjuncts and inverses.

We first report the query in the syntax used throughout the paper:

$$(?x, ?y, ?z) \leftarrow (?x, (a^- + c)^*, ?y), (?y, (b^- + d \cdot d^- \cdot e), ?z).$$

Then, we show its gMark translation into *SPARQL 1.1*:

```
PREFIX : <http://example.org/gmark/>
SELECT DISTINCT ?x ?y ?z
WHERE {  ?x (((^:p0)|(:p2)))* ?y .
         ?y ((^:p1)|(:p3/^:p3/:p4)) ?z .  }
```

Next, we show its gMark translation into *Datalog*:

```
g0(x,y) <- edge(x1, 0, x0), x = x0, y = x1.
```

```
g0(x,y) <- edge(x0, 2, x1), x = x0, y = x1.
g0(x,x) <- edge(x, r, y).
g0(y,y) <- edge(x, r, y).
g0(x,y) <- g0(x, z), g0(z, y).
g1(x,y) <- edge(x1, 1, x0), x = x0, y = x1.
g1(x,y) <- edge(x0, 3, x1), edge(x2, 3, x1),
           edge(x2, 4, x3), x = x0, y = x3.
query(x, y, z) <- g0(x, y), g1(y, z).
```

Next, we show its gMark translation into *openCypher*:

```
MATCH (x)-[:p0|p2*]->(y), (y)<-[:p1]-(z)
RETURN DISTINCT x, y, z
UNION
MATCH (x)-[:p0|p2*]->(y),
      (y)-[:p3]->()<-[:p3]-()-[:p4]->(z)
RETURN DISTINCT x, y, z;
```

Finally, we show its gMark translation into *SQL*:

```
WITH RECURSIVE c0(src, trg) AS (
    SELECT edge.src, edge.src FROM edge
    UNION
    SELECT edge.trg, edge.trg FROM edge
    UNION
    SELECT s0.src, s0.trg
    FROM (SELECT trg as src, src as trg, label
          FROM edge) as s0
    WHERE s0.label = 0
    UNION
    SELECT s0.src, s0.trg
    FROM edge s0
    WHERE s0.label = 2 )
 c1(src, trg) AS (
    SELECT src, trg FROM c0
    UNION
    SELECT head.src, tail.trg
    FROM c0 as head, c1 as tail
    WHERE head.trg = tail.src),
 c2(src, trg) AS (
    SELECT s0.src, s0.trg
    FROM (SELECT trg as src, src as trg, label
          FROM edge) as s0
    WHERE s0.label = 1
    UNION
    SELECT s0.src, s2.trg
    FROM edge s0,
         (SELECT trg as src, src as trg, label
          FROM edge) as s1, edge s2
    WHERE s0.trg = s1.src AND s1.trg = s2.src
          AND s0.label = 3 AND s1.label = 3
          AND s2.label = 4 )
SELECT DISTINCT c1.src, c2.src, c2.trg
FROM c1, c2
WHERE c2.src = c1.trg;
```